arXiv:1605.04691v1 [cs.AI] 16 May 2016

On Avoidance Learning with Partial Observability

Tom J. Ameloot*
Hasselt University and transnational University of Limburg

Abstract

We study a framework where agents have to avoid aversive signals. The agents are given only partial information, in the form of features that are projections of task states. Additionally, the agents have to cope with non-determinism, defined as unpredictability on the way that actions are executed. The goal of each agent is to define its behavior based on feature-action pairs that reliably avoid aversive signals. We study a learning algorithm, called A-learning, that exhibits fixpoint convergence, where the belief of the allowed feature-action pairs eventually becomes fixed. A-learning is parameter-free and easy to implement.

Contents

1 Introduction **1**

2 Related Work **6**

3 Fundamental Notions **7**
3.1 Tasks 7
3.2 Strategies 9

4 Avoidance Learning **11**
4.1 A-learning Algorithm 11
4.2 Results 13
4.3 Fairness 15

5 Simple Grid Navigation **16**
5.1 Definitions 16
5.2 Results 18

6 Conclusion and Further Work **20**

1 Introduction

The main aim of this paper is to let agents solve tasks by ultimately avoiding aversive signals forever. This approach entails an interesting and perhaps quite

*T.J. Ameloot is a Postdoctoral Fellow of the Research Foundation – Flanders (FWO).

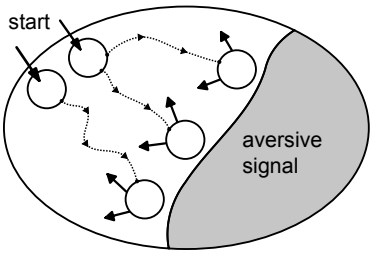

Figure 1.1: Tasks contain aversive signals that should be avoided. In general, we allow multiple start states.

strong guarantee on the agent performance. The motivation is partly to understand how animals are successful in solving problems, like navigation (Geva-Sagiv et al., 2015), with limited sensory information and unpredictable effects in the environment. The animal should find food or return home before it gets lost or becomes exhausted.

We study a general framework in which agents need to avoid problems in tasks. If the agent encounters a problem, an aversive signal is received. This way the agent could learn to avoid the problem, by avoiding the usage of actions and action-sequences that lead to aversive signals. The general idea is sketched in Figure 1.1. Before we discuss our approach, we first briefly discuss two important ingredients of the framework, namely, partial information and non-determinism.

Partial information and non-determinism First, we assume that the agent is given only partial information, as follows: each encountered task state is projected to a set of features. This is a propositional representation, where each feature is a true/false question posed about the state (van Otterlo, 2009). The number of features determines the granularity by which states can be perceived by agents. Therefore, the behavior of the agent will be based on feature-action associations, and not on (direct) state-action associations.[1] Each application can choose its own features and its own way of computing them. Examples of features are: detected edges in images, impulses through sensors, temporal events over streams, AND-OR combinations thereof, etc. In this paper, we assume that tasks have only a finite number of features, although there could still be many features. Perhaps not surprisingly, theoretical investigations show how hard it is to solve tasks under partial information, see e.g. (Lusena et al., 2001; Roy et al., 2005; Chatterjee and Chmelík, 2015).[2]

Second, we allow tasks to be non-deterministic. This means that the effect of some action-applications to states can not be predicted. In this paper, we assume that non-determinism is an inherent property of tasks. Although partial information also limits the reasoning within the agent, and therefore generally prevents accurate predictions, it remains a separate assumption to allow tasks themselves to be non-deterministic. For example, one may consider tasks in

[1] If the same feature is used by several states, this may be seen as a form of generalization over those states. In this paper, features are used directly, and we do not perform a (second) generalization step over the observed features.

[2] Part of the motivation for this paper is also to reason about the feature design for solving tasks. An example is given in Section 5.

which features actually provide complete information, and where the agent could still struggle with non-determinism.[3]

Strategies The focus of this paper to understand agents based on their behavior in tasks, which could be a useful way to understand intelligence in general (Pfeifer and Scheier, 1999). As remarked earlier, in this paper, agent behavior will be based on feature-action associations. Conceptually, we may think of the agent as having a set P of possibly allowed feature-action pairs, and whenever the agent encounters a task state s, the agent (thinks it) is allowed to perform all actions a for which there is a feature f observed in state s such that $(f, a) \in P$. We also refer to P as a policy.

We say that a set P of feature-action pairs constitutes a *strategy* for a start state if P will never lead to an aversive signal when starting from that start state. We note that it is not always sufficient for the states near the aversive signals to steer away from them, because sometimes the agent may get trapped in a zone of the state space that does not immediately give aversive signals but from which it is impossible to reliably escape the aversive signals. The agent should avoid such zones, which could require that the agent anticipates aversive signals from very early on.

Our aim in this paper is to reason about the existence of such successful strategies for classes of tasks, and to discuss an algorithm to find such strategies automatically. A main challenge throughout this study is posed by the compression of state information into features and the uncontrollable outcomes due to non-determinism.

Reward-based value estimation seems unsuitable Before presenting more details of our algorithm, we first argue that algorithms based on (numerical) reward-based value estimation do not always appear suitable for reliably finding problem-avoiding strategies.

On the theory side, convergence proofs of value estimation algorithms often require the learning step-size to decrease over time, see e.g. (Jaakkola et al., 1994; Watkins and Dayan, 1992). Intuitively, convergence of the estimated values arises because the decreasing learning step-size makes it harder and harder for the agent to learn as time progresses. However, we would like to avoid putting such limits on the agent, because: (1) it is useful to also study more flexible agents because they might sometimes better describe real-world agents; (2) in practice it might be difficult to estimate in what exact way the learning step-size should decrease; and, (3) also in practice, there are no guarantees on what the estimated values will eventually be after a certain amount of time has passed, because the estimates depend strongly on random fluctuations during task performance (due to non-determinism).

In practice, a non-decreasing step-size, although potentially useful to model flexible agents that keep learning from their latest experiences (Sutton and Barto, 1998), can lead to problems of its own. We illustrate this with the example task shown in Figure 1.2a. There is a start state 1, and two actions a and b that lead back to state 1. We assume complete information for now, i.e., state 1 is presented completely to the agent as a feature with the same information, namely, the identity of state 1. Suppose that the state-action pair

[3] We will see an example of this situation a bit later in the Introduction.

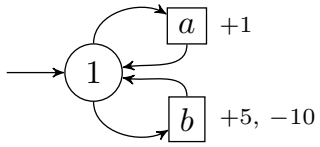

(a) Example task, with state 1, and actions a and b. Feedback to the agent is modeled as numerical reward.

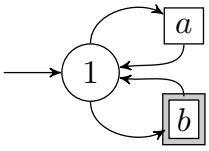

(b) The same task as in Figure 1.2a, but now feedback is modeled as (boolean) aversive signals. Applying action b results in an aversive signal, indicated by a box with double border.

Figure 1.2: An example task, represented in two ways. States and action applications are represented by circles and boxes respectively. Start states are indicated by an arrow without origin.

$(1, a)$ always gives reward $+1$. But for the pair $(1, b)$, the reward could be either $+5$ or -10. Although the pair $(1, b)$ is clearly preferable over the pair $(1, a)$ in case of positive reward, there is the risk of incurring a strong negative reward. The negative reward represents an aversive signal. In the perspective of strategies from above, note that $(1, a)$ constitutes a strategy: constantly executing action a in state 1 leads to an avoidance of aversive signals forever.

But the agent will not necessarily learn to avoid $(1, b)$ when a hidden task mechanism could periodically deceive the agent by issuing higher rewards under action b. Concretely, let n be a strictly positive natural number. To represent the outcome of action b, suppose that we constantly give reward $+5$ during the first n times $(1, b)$ is applied; the next n times we give -10; the following n times we again give $+5$, and so on. We call this the n-swap semantics. For each outcome, the empirical probability would be 0.5: indeed, the observed frequency of each outcome converges to 0.5 as we perform more applications of action b. We can choose n arbitrarily large; this does not change the empirical probability of each outcome.

Without the restriction on learning step-size, it seems that value estimation algorithms can get into trouble on the above setting because we can set n so large that after a while the agent starts to believe that the outcome would remain fixed. For example, we could start with reward $+5$ for the pair $(1, b)$ during the first n applications, and the agent starts believing that the reward really is $+5$. Then come the next n applications, where we repeatedly give reward -10, and the agent starts believing that the reward really is -10. We can swap the two outcomes forever, each for a period of n applications, and the agent will never make up its mind about the behavior of action b in state 1. This effect is illustrated in Figure 1.3.

Although the above example is very simple, real-world tasks could still exhibit problems similar to the n-swap semantics. Even if such problems are identified and understood, perhaps there are no good solutions for them as the problems might be outside the range of control for the agent. In this paper we would like to learn to avoid the aversive signals forever, even under quite adversary semantics of tasks like the n-swap semantics.

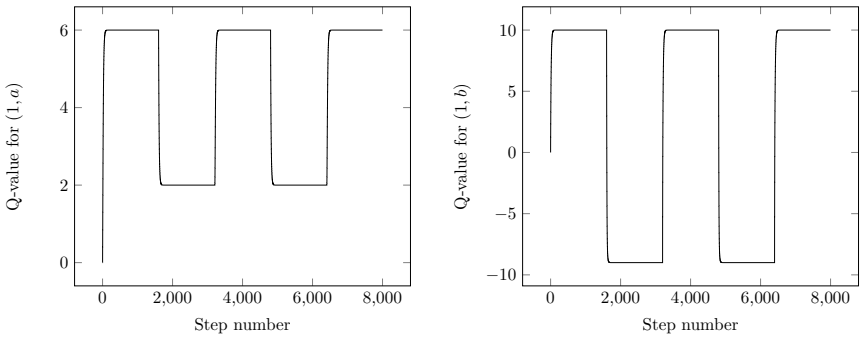

(a) Q-values for pair $(1, a)$.

(b) Q-values for pair $(1, b)$.

Figure 1.3: We have simulated the Q-learning algorithm (Watkins, 1989; Watkins and Dayan, 1992) on the example task shown in Figure 1.2a, by alternating actions a and b, but with constant step-size $\alpha = 0.25$ (and discounting factor $\gamma = 0.5$). The resulting value estimates for the state-action pairs $(1, a)$ and $(1, b)$ are plotted against time. The outcome of $(1, b)$ was either $+5$ or -10, as determined by the n-swap semantics where $n = 800$; this semantics is relative to the applications of b, and not relative to the global time steps.

Avoidance learning In the example of Figure 1.2a, we would like the agent to make up its mind more quickly that action b leads to aversive signals. An idea is to let the agent (monotonically) increase its estimate of the value of a feature-action pair. We should immediately observe, however, that this idea will not work when feedback remains to be modeled as reward, as in the example: once the outcome of $(1, b)$ is observed to be $+5$; then remembering $+5$ would lead to a preference of $(1, b)$ over $(1, a)$, causing a reward-seeking agent to (accidentally) encounter negative rewards, i.e., aversive signals, indefinitely under n-swap semantics.

Fortunately, the idea of increasing estimates seems to work when feedback is modeled with aversive signals, even in face of non-determinism. Indeed, Heger (1994) has previously proposed a learning algorithm in tasks where actions have numeric costs, representing aversive signals. By repeatedly remembering the highest observed cost for a state-action pair (with the max-operator), and by choosing actions to minimize such costs, the agent learns to steer away from high costs. We would like to further elaborate this idea and how it relates to the notion of aversion-avoiding strategies mentioned above.

In our framework, we only explicitly model aversive signals, as boolean flags: the flag "true" would mean that an aversive signal is present. This leads to a framework that is conceptually neat and computationally efficient. Because a policy is either successful in avoiding aversive signals forever, or it is not, the choice of a boolean model aligns well with our motivation to study the relationship between learning and successful strategies. To illustrate, the example of Figure 1.2a would be represented by Figure 1.2b, where only the aversive signal is explicitly represented. In general, the boolean flags will act like borders, to

demarcate undesirable areas in the state space. Reward is now only implicit: by using a strategy, as mentioned earlier, the agent can stay away from the aversive signals forever.

In the above setting with explicit aversive signals, we describe an avoidance learning algorithm, called A-learning, in which the agent repeatedly flags feature-action pairs that lead to aversive signals, or, as an effect thereof, to states for which all proposed actions are flagged (based on the observed features). Intuitively, the flags indicate "danger". On the example of Figure 1.2b, A-learning flags $(1, b)$ at the first occurrence of an aversive signal under action b; and, importantly, the strategy $(1, a)$ is never flagged.[4] There is no second chance for changing the agents mind. This gives one of the strongest convergence notions in learning, namely, fixpoint convergence, where the agent eventually stops changing its mind about the outcome of actions.

If there really is a strategy, avoidance learning will carve out a subset of good feature-action pairs from the mass of all feature-action pairs. This way, it seems that avoidance learning could be useful in making the agent eventually avoid aversive signals forever. This provides the guaranteed agent performance we would like to better understand, as remarked at the beginning of the Introduction.

Meaning of optimality In this paper we view an agent as being optimal if it can (learn to) avoid aversive signals forever. There is no explicit concept of reward. Depending on the setting, or application, aversive signals can originate from diverse sources and together they can describe a very detailed image of what the agent is allowed to do, and what the agent is not allowed to do. One obtains a rich conceptual setting for reasoning about agent performance.

For example, suppose a robotic agent should learn to move boxes in a storehouse as fast as possible. We could emit an aversive signal when the robot goes beyond a (reasonable) time limit. Any other constraints, perhaps regarding battery usage, can be combined with the first constraint by adding more signals.

Outline This paper is organized as follows. We discuss related work in Section 2. We introduce fundamental concepts like tasks, and strategies, in Section 3. We present and analyze our avoidance learning algorithm in Section 4. One of our results is that if there is a strategy for a start state then the algorithm will preserve the strategy. This mechanism can be used to materialize strategies if they exist. To better understand the nature of strategies, we prove the existence of strategies for a family of grid navigation tasks in Section 5.

2 Related Work

The idea of avoiding aversive signals, or problems in general, is related to safe reinforcement learning (García and Fernández, 2015). There, the goal is essentially to perform reinforcement learning, often based on approximation techniques for optimizing numerical reward, with the addition of avoiding certain

[4] In our description of A-learning (Section 4), the flagged feature-action pairs are removed from the agent's memory.

problematic areas in the task state space. An example could be to train a robot for navigation tasks but while avoiding damage to the robot as much as possible. In the current paper, feedback to the agent consists of the aversive signals. Reward becomes more implicit, as it lies in the avoidance of aversive signals. Therefore, the viewpoint in this paper is that the agent is called optimal when it eventually succeeds in avoiding all aversive signals forever; there is no notion of optimizing reward. The approach is related to a trend identified by García and Fernández (2015), namely, the modification of the optimality criterion.

The work by Heger (1994) is closely related to our work. The framework by Heger (1994) provides feedback to the agent in the form of numerical cost signals, which, from the perspective of this paper, could be seen as aversive signals. Similar to our n-swapping example in the Introduction (Figure 1.2a), Heger (1994) provides other examples to motivate that estimation of expected values is not suitable for reliably deciding actions. The learning algorithm proposed by Heger (1994) maps each state-action pair to the worst outcome (or cost), by means of the max-operator. By remembering the highest incurred cost for a state-action pair, the agent in some sense learns about "walls" in the state space that constrain its actions towards lower costs. The avoidance learning algorithm discussed in this paper (Section 4) is similar in spirit to the one by Heger (1994). A deviation, however, is that we assume here a boolean interpretation of aversive signals, which leads to a neat and computationally efficient framework. We additionally identify the concept of strategies, under which the agent can avoid aversive signals forever. Our interest lies in understanding such avoidance strategies and their relationship to the avoidance learning algorithm. Moreover, we also focus on partial information, by letting the agent only observe features instead of full states.

3 Fundamental Notions

3.1 Tasks

For a set X, let $\mathcal{P}(X)$ denote the powerset of X, i.e., the set of all subsets of X. A task is a tuple $T = (S, B, A, F, \delta, \varphi, \Omega)$ where

- S is a nonempty set of states;
- $B \subseteq S$ is a finite subset of start states;[5]
- A is a nonempty finite set of actions;
- F is a nonempty finite set of features;
- $\delta : S \times A \to \mathcal{P}(S)$ is the transition function;
- $\varphi : S \to \mathcal{P}(F)$ is the feature function; and
- $\Omega \subseteq S \times A$ is the set of aversive signals,

where all states $s \in S$ are reachable in the sense that there is a sequence $s_0, a_0, s_1, a_1, \ldots, s_n$ with $s_0 \in B$, $s_n = s$, and $s_i \in \delta(s_{i-1}, a_{i-1})$ for each $i \in \{1 \,..\, n\}$.

[5]The symbol B stands for "begin".

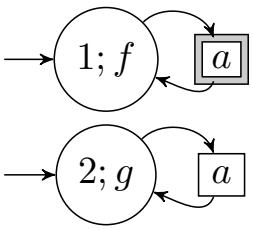

Figure 3.1: The task from Example 3.1. The basic graphical notation is explained in Figure 1.2. Inside the circles, we write the state identifier followed by a semicolon and the features of the state.

The function δ maps each pair $(s, a) \in S \times A$ to a set of possible successor states, representing non-determinism. The function φ associates a set of features to each state; an agent interacting with the task can only observe states through features and can therefore not directly observe states. The meaning of a pair $(s, a) \in \Omega$ is that the agent could witness an aversive signal when performing action a in state s.[6]

Example 3.1. We define an example task $T = (S, B, A, F, \delta, \varphi, \Omega)$ as follows: $S = \{1, 2\}$; $B = S$; $A = \{a\}$; $F = \{f, g\}$; regarding δ, we define

$$\delta(1, a) = \{1\},$$
$$\delta(2, a) = \{2\};$$

regarding φ, we define

$$\varphi(1) = \{f\},$$
$$\varphi(2) = \{g\};$$

and, we define $\Omega = \{(1, a)\}$. The task is depicted in Figure 3.1. □

Remark 3.2 (All features, one successor state)**.** Note that the function φ maps each state to a set of features. Similarly, the function δ maps each state-action pair to a set of successor states. However, an agent interacts with each function in a different way, as follows. For a state s, we assume that an agent can always observe all features in $\varphi(s)$ simultaneously. This way, the function φ may be viewed as being deterministic. In contrast, for a state-action pair (s, a), we select only one successor state from $\delta(s, a)$ to proceed with the task.

The function φ remains deterministic throughout this paper. The framework still allows us to consider tasks in which the agent can sometimes observe a certain feature and sometimes it can not. Thereto we can define richer states, in which, say, the status of sensors is stored; if a state s says that a sensor is broken, then $\varphi(s)$ could omit the feature that would otherwise be generated by the sensor. □

[6] In a fair task, if the agent would infinitely often perform action a in state s, then the agent witnesses an aversive signal infinitely often during the application of a at state s, but this signal could sometimes be omitted. See also Section 4.3.

Remark 3.3 (Modeling flexibility)**.** Our definition of task resembles that of a standard Markov decision process (Sutton and Barto, 1998), but we have added features and aversive signals. There can be many features, actions, and start states. And we allow an infinite number of states. □

3.2 Strategies

Since the agent may only see features, and not states directly, agent behavior has to be based on feature-action associations.

Let $T = (S, B, A, F, \delta, \varphi, \Omega)$ be a task. A *policy* for T is a total function $\pi : F \to \mathcal{P}(A)$. We allow features to be mapped to empty sets of actions. If the task is understood from the context, for a state $s \in S$ we define

$$act(s, \pi) = \bigcup_{f \in \varphi(s)} \pi(f),$$

i.e., $act(s, \pi)$ is the set of all actions that are proposed by the policy π based on the features in s. We say that a state s is *blocked in* π if $act(s, \pi) = \emptyset$, i.e., the policy does not propose actions for s.

Remark 3.4 (Features as actors)**.** For a state s, we do not view $\varphi(s)$ as an atomic signature to which actions should be associated. Instead, the definition of $act(s, \pi)$ indicates that each feature in $\varphi(s)$ may independently propose its own actions, regardless of what is proposed by other features. All proposed actions are collected into a set, by means of the union-operator.[7] Therefore, features are little actors that become active at appropriate times and that suggest to the agent what actions are (supposedly) allowed. This viewpoint resembles the way that an individual neuron (or a small group of neurons) in the brain could represent a distinct concept and could be individually linked to actions (Potjans et al., 2011; Frémaux et al., 2013).[8] It is the goal of the learning algorithm (Section 4) to remove feature-action associations that lead to aversive signals or, as a result of such removals, to blocked states. □

We now consider the following definition:

Definition 3.5. A policy π is called a *strategy* for a start state $s_0 \in B$ if

1. $act(s_0, \pi) \neq \emptyset$;
2. $\forall s \in S$, $\forall a \in act(s, \pi)$,
   (a) $\forall s' \in \delta(s, a)$ we have $act(s', \pi) \neq \emptyset$; and,
   (b) $(s, a) \notin \Omega$.

□

[7] It would in general not be possible to replace this with an intersection-operator, because then there might be some unreliable features to which no action can be associated. Under an intersection-operator, such empty sets will destroy the proposals contributed by reliable features.

[8] So, features constitute the agent's mind, like neurons constitute an organism's brain. Likely there are only finitely many neurons, so the assumption of having a finite number of features and a finite number of actions (or decisions) appears natural.

In words: a policy is a strategy for a start state s_0 if the policy acts upon s_0; and, for any states upon which the policy acts, the reached successor states can also be acted upon, and the policy never causes aversive signals. Intuitively, to use a strategy, for each encountered state s we first select some (arbitrary) feature $f \in \varphi(s)$ that satisfies $\pi(f) \neq \emptyset$, and we subsequently select an arbitrary action $a \in \pi(f)$.

Remark 3.6 (Global viewpoint)**.** The definition of strategy demands properties in a global fashion, possibly also for states that would not be explored when strictly following the strategy. This condition however ensures that learning algorithms can never have negative experiences when they perform actions suggested by the strategy; see Section 4.

Suppose π is a strategy, and let f be a feature with $\pi(f) \neq \emptyset$. Intuitively, the definition of strategy says that f is a reliable feature, in the sense that every time we see it, we may safely perform all actions in $\pi(f)$, without the risk of encountering blocked states and aversive signals. This is related to the Markov assumption (Sutton and Barto, 1998), because we do not have to remember any features that were seen during previous time steps, and we may instead choose actions based on just f by itself. □

Example 3.7. Consider the task from Example 3.1. There is no strategy for start state 1, but there is a strategy π for start state 2 defined as: $\pi(f) = \emptyset$ and $\pi(g) = \{a\}$. □

The following property illustrates that strategies are resilient to adding new features. In practical applications, this means that the addition of new kinds of features will not destroy previously existing strategies.[9]

Proposition 3.8. Let $T_1 = (S_1, B_1, A_1, F_1, \delta_1, \varphi_1, \Omega_1)$ be a task. Let V be a set of features that is disjoint from F_1. Let $T_2 = (S_2, B_2, A_2, F_2, \delta_2, \varphi_2, \Omega_2)$ be another task that is almost the same as T_1 except that $F_2 = F_1 \cup V$ and for each state s the constraint $\varphi_2(s) \cap F_1 = \varphi_1(s)$ holds.[10] Let s_0 be a start state, and suppose that a policy $\pi : F_1 \to \mathcal{P}(A_1)$ is a strategy for s_0 in T_1. Then π is also a strategy for s_0 in T_2.

Proof. We show that the conditions of strategy in Definition 3.5 are satisfied for π in T_2. To better show which task is involved, for a state s and a task index $i \in \{1, 2\}$, we write

$$act_i(s, \pi) = \bigcup_{f \in \varphi_i(s)} \pi(f).$$

When π is used in T_2, we assume $\pi(f) = \emptyset$ for each $f \in V$. Above we have also assumed that $A_1 = A_2$.

Condition 1 Since π is a strategy for s_0 in T_1, we have $act_1(s_0, \pi) \neq \emptyset$. This implies there is some $f \in \varphi_1(s_0)$ with $\pi(f) \neq \emptyset$. Since $\varphi_1(s_0) \subseteq \varphi_2(s_0)$ by assumption, we obtain $act_2(s_0, \pi) \neq \emptyset$.

[9] For example, in a robotics application, new features could be the result of adding new sensor types to the robot.

[10] This means that T_2 uses the features of F_1 in the same way as T_1.

Condition 2 Let s be a state and assume there is some action $a \in act_2(s, \pi)$. First we argue that $a \in act_1(s, \pi)$. There must be a feature $f \in \varphi_2(s)$ with $a \in \pi(f)$. But since π only knows features in F_1, we have $f \in \varphi_2(s) \cap F_1 = \varphi_1(s)$. Hence, $a \in act_1(s, \pi)$.

We first handle Condition 2a. Let $s' \in \delta(s, a)$. Because $a \in act_1(s, \pi)$ and π satisfies Condition 2a in T_1, we know $act_1(s', \pi) \neq \emptyset$. So there is a feature $f' \in \varphi_1(s')$ with $\pi(f') \neq \emptyset$. Since $\varphi_1(s') \subseteq \varphi_2(s')$, we know $a \in act_2(s', \pi) \neq \emptyset$, as desired.

Now we handle Condition 2b. Because $a \in act_1(s, \pi)$ and π satisfies Condition 2b in T_1, we know $(s, a) \notin \Omega_1$. Since $\Omega_2 = \Omega_1$, we obtain $(s, a) \notin \Omega_2$, as desired. □

4 Avoidance Learning

We present and study an avoidance learning algorithm, and its relationship to the concept of strategy introduced in Section 3.2.

4.1 A-learning Algorithm

Algorithm 4.1 is an avoidance learning algorithm. The algorithm describes how the agent interacts with the task, and how feature-action combinations are forgotten as the direct or indirect result of aversive signals. Some aspects of the interaction are not under control of the agent, in particular how a successor state is chosen by means of function δ, and how features are derived from states by means of function φ. We now provide more discussion of the algorithm. Henceforth, we will refer to Algorithm 4.1 as A-learning.

The essential product of A-learning is a set $P \subseteq F \times A$ that represents the allowed feature-action pairs; the symbol P stands for "possibilities". At any time, the set P uniquely defines a policy π as follows: for each $f \in F$, we define $\pi(f) = \{a \in A \mid (f, a) \in P\}$. Regarding notation, for any state s, we write $act(s, P)$ to denote the set $act(s, \pi)$ of proposed actions, where π is the unique policy defined by P.

We now explain the steps of A-learning in more detail.

- Line 1 initializes P with all feature-action pairs. We will gradually remove pairs if they lead to Ω or to blocked states (that are created by removals of the first kind).
- Line 2 selects a random start state. The control flow of the algorithm is redirected here each time we want to restart the task. But we never re-initialize P.

  Task restarts may be requested by A-learning itself (see below), or externally by the training framework in which A-learning is running.
- Line 3 requests a task restart in case the chosen start state is blocked. This allows more exploration from the other start states. As we will see later in Theorem 4.5(1), if no actions remain for a start state then this start state has no strategy.
- At Line 6, the algorithm enters a learning loop. The loop is only exited to satisfy task restart requests, at Line 7.

Algorithm 4.1: Avoidance learning (A-learning)

Data: Task $T = (S, B, A, F, \delta, \varphi, \Omega)$.

```
1  P := F × A;
2  s := choose from B;
3  if act(s, P) = ∅ then
4  |  request restart (see below);
5  end
6  while true do
7  |  if restart requested then
8  |  |  go to Line 2;
9  |  end
10 |  a := choose from act(s, P);
11 |  s' := choose from δ(s, a);
12 |  if (s, a) ∈ Ω or act(s', P) = ∅ then
13 |  |  P := P \ (φ(s) × {a});
14 |  |  request restart;
15 |  end
16 |  s := s';
17 end
```

- At Line 10, we choose an action a to apply to current state s based on the set $act(s, P)$ of still allowed actions. At Line 11, we are subsequently given a successor state s', chosen arbitrarily from $\delta(s, a)$.
- Next, at Line 12, we check whether we have encountered Ω or if successor state s' is blocked. In either case we exclude from P the feature-action pairs that caused us to apply action a in state s (Line 13), and we restart the task (Line 14).
- If we do not encounter Ω and state s' is not blocked, then we proceed with the while loop (Line 16).

Note that in general there are multiple runs of A-learning on a task, because of the choice on action selection and the choice on successor state. Each run of A-learning is infinitely long. Nonetheless, there is always an eventual fixpoint on the set P because after the initialization we only remove feature-action pairs. There are only a finite number of possible feature-action pairs, although there could be many. When the run is clear from the context, we write P^* to denote the fixpoint of P obtained in that run.

For conceptual convenience, we can divide each run of A-learning into trials by using the task restarts as dividers: whenever we execute Line 2, the previous trial ends and the next trial begins. Each trial is thus a sequence $s_0, a_0, s_1, a_1, \ldots, s_n$, where s_0 is a start state, s_n is the last state of the trial, and $s_i \in \delta(s_{i-1}, a_{i-1})$ for each $i \in \{1 \,..\, n\}$.[11]

[11] We only use explicit task restarts (Line 2) to divide runs into trials, and not the encounter

Remark 4.1 (No stopping condition)**.** There is no stopping condition in the algorithm because in general we may not be able to detect when the agent has explored the task sufficiently to be successful at avoiding aversive signals. □

Remark 4.2 (Greediness)**.** We would like to emphasize that A-learning is always greedy in avoiding Ω. This is an important deviation from the ϵ-greedy exploration principle (Sutton and Barto, 1998), where at each time step the agent chooses a random action with small probability $\epsilon \in [0, 1]$. We do not use that mechanism here because otherwise the agent keeps running the risk of encountering aversive signals (García and Fernández, 2015). □

Remark 4.3 (Internal task restarts)**.** The reason for requesting a task restart at Line 14 is that sometimes the agent could become stuck in a zone of the state space where there are only blocked states or aversive signals. In that case, if we want the agent to start removing feature-action pairs to prevent future aversive signals, we should first transport the agent to a zone in the state space without blocked states and aversive signals. For example, in a robot navigation problem, the robot could learn to avoid pits, but once it enters a pit it can perhaps not reliably escape without the help of an external supervisor. □

Remark 4.4 (Memory efficiency)**.** Algorithm 4.1 explicitly stores the allowed feature-action pairs in a set P. This is an intuitive perspective for the theory developed in this paper. However, in practice it may sometimes be more efficient to store the opposite information, namely, the removed feature-action pairs. This way all allowed feature-action pairs can still be uniquely recovered. Using the analogy of a planar map, where aversive signals are borders between neutral zones on the one hand and undesirable zones on the other hand, there could be a decreased memory usage in storing only the border (i.e., the removed feature-action pairs) if the borders are simple shapes instead of irregular shapes with many protrusions. □

4.2 Results

The following theorem helps to understand what A-learning computes.

Theorem 4.5. For all tasks $T = (S, B, A, F, \delta, \varphi, \Omega)$, for each $s_0 \in B$, for each run of A-learning, where P^* denotes the fixpoint,

1. if there is a strategy for s_0 then $act(s_0, P^*) \neq \emptyset$.
2. if $act(s_0, P^*) \neq \emptyset$ then every trial for s_0 after the fixpoint avoids blocked states and Ω.

Proof. We consider the two properties separately.

Property 1 Suppose there is a strategy π for s_0. We show that the feature-action pairs of π are preserved in P^*, so that $act(s_0, \pi) \neq \emptyset$ would imply $act(s_0, P^*) \neq \emptyset$. Towards a contradiction, suppose that A-learning removes a pair (f, a) from P where $a \in \pi(f)$; let (f, a) be the first such pair that is removed. The removal has happened as follows: we reach a state s with $f \in \varphi(s)$

of start states. This means that in principle we allow $s_i \in B$ for some or all $i \in \{1 \,..\, n\}$.

and we perform a, and either the successor state $s' \in \delta(s, a)$ is blocked or we receive an aversive signal. We discuss each case in turn.

Let P denote the remaining feature-action pairs just before we remove (f, a). Note that $a \in \pi(f)$ and $f \in \varphi(s)$ together imply $a \in act(s, \pi)$.

- Suppose that s' is blocked. Since π is a strategy, by condition 2a of Definition 3.5, we have assumed $act(s', \pi) \neq \emptyset$. So, there is a feature $f' \in \varphi(s')$ and an action $a' \in \pi(f')$. Since (f, a) is the first pair of π that is removed, we still have $(f', a') \in P$. But then $a' \in act(s', P)$, and s' is actually not blocked; we have found a contradiction.

- Suppose that an aversive signal was received when applying a to s, which implies $(s, a) \in \Omega$. This immediately contradicts the assumption that π satisfies condition 2b of Definition 3.5.

Property 2 Suppose $act(s_0, P^*) \neq \emptyset$. Towards a contradiction, suppose that after the fixpoint there is a trial for start state s_0 where we encounter a state s and we perform an action a such that either the successor state is blocked or we receive an aversive signal. Suppose we conceptually halt the offending trial at the first encountered problem. We have followed a path:

$$s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} \ldots s_{n-1} \xrightarrow{a_{n-1}} s_n = s \xrightarrow{a_n = a} s',$$

for some $s' \in \delta(s, a)$. We have $a_i \in act(s_i, P^*)$ for each $i \in \{0 \,..\, n\}$. We note in particular that $a \in act(s, P^*)$. Next we distinguish two cases, depending on the type of problem.

- Suppose that $act(s', P^*) = \emptyset$. Then A-learning now removes $\varphi(s) \times \{a\}$ from P. But then we will no longer propose action a for state s, which was previously allowed by the fixpoint. Then P^* would be an invalid fixpoint, which is a contradiction.

- Suppose that an aversive signal is received when applying a to s, which implies $(s, a) \in \Omega$. We make a similar reasoning as in the previous case: A-learning removes $\varphi(s) \times \{a\}$ from P. Again P^* would be an invalid fixpoint.

□

Remark 4.6 (Strategies and eventual success)**.** Suppose that a task has a strategy for each start state. In that case, Theorem 4.5 tells us that every run of A-learning will eventually avoid blocked states and aversive signals. The agent therefore makes a transition from first discovering the strategies to later exploiting the strategies.

The opposite is not necessarily true: there are tasks for which exist runs that eventually avoid blocked states and aversive signals, but without there being a strategy in the sense of Definition 3.5. This is illustrated by the task in Figure 4.1. Consider a run where the first application of action b in state 1 results in an aversive signal, and after which we immediately restart the task. In that run, there is no further exploration to state 2, which causes $(f, a) \in P^*$; hence, $act(1, P^*) = \{a\} \neq \emptyset$. However, note that if the internal restart request at Line 14 of Algorithm 4.1 would sometimes not be handled immediately, but a few steps later, then some runs will not preserve the pair (f, a). □

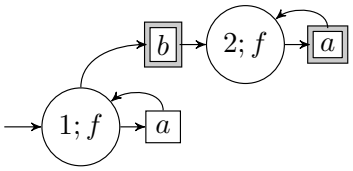

Figure 4.1: An example task that has no strategy for the start state 1. But there is a run of A-learning in which the feature-action pair (f, a) is preserved. The graphical notation is explained in Figure 3.1.

Remark 4.7 (Usage of A-learning)**.** The insights of Theorem 4.5 could be used as follows. First, although proving that a strategy exists helps in understanding guarantees on the agent performance, programming the strategy by hand could be tedious and time-consuming. So, Property 1 could be used to materialize strategies once they are proven to exist.

Second, if one does not know whether a strategy exists, Property 2 could be used to perform a preliminary search for strategies. Although the discovered strategies might not be easily interpreted, they could serve as inspiration for a theoretical study of strategies for the tasks at hand. A practical consideration, however, is that it might not be possible to efficiently detect the fixpoint, i.e., typically one does not know if a fixpoint has been reached when A-learning has not removed feature-action pairs for a while. □

4.3 Fairness

So far we have silently allowed all possible runs of A-learning. For example, we did not explicitly demand that the agent actually must receive an aversive signal when applying an action a to a state s where $(s, a) \in \Omega$. The aversive signal could also be omitted. This brings us to the topic of fairness (Francez, 1986).

Intuitively, for this paper, fairness would mean that there is sufficient exploration of the task. A practical application of A-learning (Algorithm 4.1) could take the following fairness assumptions into account:

- if we execute Line 2 infinitely often then we choose each start state infinitely often;
- to fully learn the task from each start state, we infinitely often issue external task restarts at Line 7; those restarts are not requested by A-learning itself;
- at Line 10, if we encounter the same pair of a state s and set P infinitely often then we choose each action $a \in act(s, P)$ infinitely often;
- at Line 11, if we apply action a infinitely often to state s then each successor state $s' \in \delta(s, a)$ is visited infinitely often from an application of a to s;

- at Line 12, if we perform action a in state s infinitely often, where $(s, a) \in \Omega$, then the agent should infinitely often receive an aversive signal when applying a to s;

The only aspect of fairness that can be directly influenced by the agent itself, is the action selection at Line 10. For this purpose, a random number generator can be used to select random indices in an array-representation of the proposed actions.[12]

Remark 4.8 (Fairness not required)**.** Note that Theorem 4.5 also works for unfair runs. Every run has a fixpoint on P, whether the run is fair or not. But by exploring fewer states, or by issuing fewer aversive signals, an unfair run essentially makes it easier for the agent to avoid aversive signals. This way, some feature-action pairs could remain forever, even though a more fair exploration of the task could have removed them.

Also, because the notion of strategy in Definition 3.5 is rather strong, it is not possible for a fair run or an unfair run to confront the agent with a situation that leads to the failure of a strategy. The agent will never be disappointed in the exploitation of the strategy. □

5 Simple Grid Navigation

We study a simple class of grid navigation problems.

5.1 Definitions

Let $\mathbb{Z}$ denote the set of integers. For any two points $p_1, p_2 \in \mathbb{Z} \times \mathbb{Z}$, denoting $p_1 = (x_1, y_1)$ and $p_2 = (x_2, y_2)$, we recall the definition of L_1-distance between p_1 and p_2:

$$d_1(p_1, p_2) = abs\,(x_2 - x_1) + abs\,(y_2 - y_1)\,.$$

A *simple grid navigation problem* is a quintuple $\mathcal{G} = (\mathit{Width}, \mathit{Height}, \mathit{Starts}, \mathit{Targets}, \tau)$, where

- $\mathit{Width} \in \mathbb{N}$ and $\mathit{Height} \in \mathbb{N}$ are the dimensions of a terrain;
- $\mathit{Starts} \subseteq \{0\,..\,\mathit{Width}\} \times \{0\,..\,\mathit{Height}\}$ is a set of start locations;
- $\mathit{Targets} \subseteq \{0\,..\,\mathit{Width}\} \times \{0\,..\,\mathit{Height}\}$ is a set of possible target locations; and,
- $\tau \in \mathbb{N}$ is a time limit,

with the following assumptions,

- $\forall p, q \in \mathit{Targets}$, we assume $d_1(p, q) < \tau$; and,
- $\forall p \in \mathit{Starts}, \forall q \in \mathit{Targets}$, we assume $d_1(p, q) < \tau$.

[12] One has to assume that the random number generator is fair in selecting all indices infinitely often if we let the system run forever.

The intuition is that at the beginning of a session we select a start location $p \in \mathit{Starts}$ and an initial active target location $q \in \mathit{Targets}$ and we should navigate from p to q within time τ. Whenever we reach the active target location q we choose another target location $q' \in \mathit{Targets}$ and we should now navigate from q to q' within time τ. This relocation of the active target may be repeated an arbitrary number of times. But at any moment we may also begin a new session, in which we again choose a start location and initial target location. There are infinitely many sessions. The available actions are: left, right, up, down, left-up, left-down, right-up, right-down, and wait. Importantly: failure to respect the time τ results in an aversive signal; we aim to eventually avoid such aversive signals.

For a location $(x, y) \in \mathbb{N} \times \mathbb{N}$ and an action a, we now define the possible successor locations that result from the application of a to (x, y); we denote this set as $\mathit{move}(x, y, a)$. A set of multiple possible successors is used to represent non-determinism. An empty set of of successors is used to say that the action would lead outside the considered terrain. We assume the following actions to be deterministic: left, right, up, and down. The other, "diagonal", actions are non-deterministic. For example, for each $(x, y) \in \mathbb{Z} \times \mathbb{Z}$,

$$\mathit{move}(x, y, \text{left}) = \{(x - 1, y)\},$$

$$\mathit{move}(x, y, \text{left-up}) = \{(x - 1, y - 1), (x - 1, y), (x, y - 1)\},$$

and,

$$\mathit{move}(x, y, \text{wait}) = \{(x, y)\}.$$

We make the assumption that the direction of the positive Y-axis corresponds to "downward".

We now define the task structure $\mathit{task}(\mathcal{G}) = (S, B, A, F, \delta, \varphi, \Omega)$ that corresponds to the above grid problem $\mathcal{G}$. Here it will be convenient to view states and features as structured objects, with components; for an object x with a component y, we write $x.\text{y}$ to access the component.

- the set S consists of all triples s with components *agent*, *target*, and *time*, satisfying the following constraints: $s.\text{agent}$ and $s.\text{target}$ are both in the set $\{0 \,..\, \mathit{Width}\} \times \{0 \,..\, \mathit{Height}\}$, and $s.\text{time} \in \{0 \,..\, \tau\}$;
- the set B consists of those states s where $s.\text{agent} \in \mathit{Starts}$, $s.\text{target} \in \mathit{Targets}$, and $s.\text{time} = \tau$;
- $A = \{\text{left, right, up, down, left-up, left-down, right-up, right-down, wait}\}$;
- the set F consists of all pairs f with components *offset* and *time*, satisfying the constraints: $f.\text{offset} \in \{-\mathit{Width} \,..\, \mathit{Width}\} \times \{-\mathit{Height} \,..\, \mathit{Height}\}$ and $f.\text{time} \in \{0 \,..\, \tau\}$;[13]
- the transition function δ is described by Algorithm 5.1; for a state $s \in S$ and action $a \in A$, the set $\delta(s, a)$ consists of all states that could possibly be returned by Algorithm 5.1 upon receiving input (s, a);

[13] These features represent an agent-centric perspective, in which the relative offset to the target is stored (see below).

- regarding φ, for each $s \in S$, we define $\varphi(s) = \{f\}$ where $f \in F$ is the single feature for which $f.\text{offset} = s.\text{target} - s.\text{agent}$ and $f.\text{time} = s.\text{time}$; and,
- $\Omega = \{(s, a) \in S \times A \mid s.\text{time} = 0\}$.

Algorithm 5.1: Action application for grid navigation (Section 5)

input : (1) current state $s \in S$
(2) action $a \in A$

output: successor state $s' \in S$

$(x_1, y_1) := s.\text{agent}$;
$(x_2, y_2) :=$ choose from $move(x_1, y_1, a)$;
if $(x_2, y_2) \in \{0\,..\,Width\} \times \{0\,..\,Height\}$ **then**
  $s'.\text{agent} := (x_2, y_2)$;
else
  $s'.\text{agent} := (x_1, y_1)$;
end
if $s'.agent = s.target$ **then**
  $s'.\text{target} :=$ choose from $Targets$;
  $s'.\text{time} := \tau$;
else
  $s'.\text{target} := s.\text{target}$;
  $s'.\text{time} := \max(0, s.\text{time} - 1)$;
end

5.2 Results

Proposition 5.1. For each grid problem $\mathcal{G}$, there is a strategy for each start state in $task(\mathcal{G})$.

Proof. Denote $task(\mathcal{G}) = (S, B, A, F, \delta, \varphi, \Omega)$. We define one policy that is a strategy for all start states.

First, we define an auxiliary set $V \subseteq F$ to consist of all features f for which

$$|f.\text{offset}|_1 < f.\text{time},$$

where $|(x, y)|_1 = abs\,(x) + abs\,(y)$ is the L_1-norm of a point (x, y). Intuitively, such features indicate that the deterministic distance from the agent location to the target location – where we only use the actions left, right, up, and down – can be bridged within the remaining time.

We now define a policy π. For all $f \in F \setminus V$ we define $\pi(f) = \emptyset$, and for each $f \in V$, denoting $f.\text{offset} = (x, y)$, we define

$$\pi(f) = \begin{cases} \{\text{left}\} & \text{if } x < 0; \\ \{\text{right}\} & \text{if } x > 0; \\ \{\text{up}\} & \text{if } x = 0 \text{ and } y < 0; \\ \{\text{down}\} & \text{if } x = 0 \text{ and } y > 0; \\ \{\text{wait}\} & x = 0 \text{ and } y = 0. \end{cases}$$

As mentioned earlier, we define downwards as the direction of the positive Y-axis. The case where $\pi(f) = \{\text{wait}\}$ occurs when the agent is located at the target.[14]

Let $s_0 \in B$. We show that π is a strategy for s_0, according to Definition 3.5.

Condition 1 of Definition 3.5 We show that $act(s_0, \pi) \neq \emptyset$. By assumption on s_0, we have $s_0.\text{agent} \in \mathit{Starts}$, $s_0.\text{target} \in \mathit{Targets}$, and $s_0.\text{time} = \tau$. By using the distance assumptions on locations in $\mathcal{G}$, we obtain $d_1(s_0.\text{agent}, s_0.\text{target}) < \tau$. Letting f be the single feature in $\varphi(s_0)$, we see that $|f.\text{offset}|_1 < \tau = f.\text{time}$, which implies that $f \in V$. Hence $\pi(f) \neq \emptyset$, which implies $act(s_0, \pi) \neq \emptyset$.

Condition 2a of Definition 3.5 Let $s \in S$. Suppose there is some action $a \in act(s, \pi)$. Let f denote the single feature of s. We have $a \in \pi(f)$, which implies $f \in V$.

Let $s' \in \delta(s, a)$. We must show that $act(s', \pi) \neq \emptyset$. Let f' be the single feature of s'. We will show that $f' \in V$, which implies $\pi(f') \neq \emptyset$, and further that $act(s', \pi) \neq \emptyset$. Based on Algorithm 5.1, we reason about what has happened during the application of action a to state s.

- Suppose the if-test at Line 8 succeeds, i.e., the agent reaches the target location. Then $s'.\text{time} = \tau$, and

  $$\begin{aligned} d_1(s'.\text{agent}, s'.\text{target}) &= d_1(s.\text{target}, s'.\text{target}) \\ &< \tau, \end{aligned}$$

  where we use the distance assumption between target locations. Overall, $|f'.\text{offset}|_1 < f'.\text{time}$; hence, $f' \in V$.

- Suppose the if-test at Line 8 does not succeed, i.e., the agent did not yet reach the target location. It must be that $a \neq \text{wait}$, because otherwise $f.\text{offset} = (0, 0)$, which implies $s'.\text{agent} = s.\text{agent} = s.\text{target}$, and the test at Line 8 would have succeeded (see previous case). So, $a \in \{\text{left}, \text{right}, \text{up}, \text{down}\}$.

  First, we observe that

  $$d_1(s'.\text{agent}, s'.\text{target}) < d_1(s.\text{agent}, s.\text{target}).$$

  Indeed, this property holds because (1) the locations $s.\text{agent}$ and $s.\text{target} = s'.\text{target}$ are inside the convex terrain; (2) the action a is given deterministic movement semantics (i.e., there is precisely one outcome), causing $s'.\text{agent}$ to be both inside the terrain and strictly closer to $s.\text{target} = s'.\text{target}$.

  Second, we also observe that

  $$f'.\text{time} = f.\text{time} - 1,$$

  since $s'.\text{time} = \max(0, s.\text{time} - 1)$ by definition and $s.\text{time} > 0$ (which follows from $f \in V$).

[14] Algorithm 5.1 implies that the situation where $f.\text{offset} = (0, 0)$ only occurs when the agent reaches some target location and the next target location is the same as the old target location.

Overall, we may now write

$$\begin{aligned} |f'.\text{offset}|_1 &< |f.\text{offset}|_1 \\ &\leq f.\text{time} - 1 \\ &= f'.\text{time}. \end{aligned}$$

In the second line we have used $f \in V$. We conclude that $f' \in V$.

Condition 2b of Definition 3.5 Let $s \in S$. Suppose there is some action $a \in act(s, \pi)$, which implies that the single feature f of s must be in V. By definition of V, we have $|f.\text{offset}|_1 < f.\text{time}$. Hence, $s.\text{time} > 0$, which implies $(s, a) \notin \Omega$, as desired. □

Remark 5.2 (Richness in strategy)**.** The policy defined in the proof of Proposition 5.1 is in general not the maximal strategy, in the sense that the policy could be extended with more actions than currently specified. For instance, if the time limit is high then the agent can randomly wander around before it becomes sufficiently urgent to reach a target location. The agent may also use the diagonal actions, like left-up, if the time limit is not violated under either of the three outcomes. □

Remark 5.3 (Extendability)**.** It is possible to extend the above setting of grid navigation to richer state representations, by including for example the locations of additional objects (that do not influence the agent). If this new information would be communicated to the agent with a set of features that is disjoint from the set of old features in Section 5.1, then Proposition 3.8 tells us that the strategy described in the proof of Proposition 5.1 is still valid. □

6 Conclusion and Further Work

We have used the notion of strategies to reason about the successful avoidance of aversive signals in tasks. We have shown that our avoidance learning algorithm always preserves those strategies. Now we discuss some interesting topics for further work.

Feature detectors In this paper we have considered a framework in which features are essentially black boxes, in the sense that we do not assume anything about the way that they are computed. Hence, we do not know how features are related to the task environment. It would be interesting to develop more detailed insights into how features can be designed, to ensure that strategies, or similarly successful policies, are possible.

In particular, it seems fascinating to explore possible connections between our framework and neuron-like models, where features would be represented by neurons or by small groups of neurons. It is currently an open question whether or not feature learning in the brain is a completely unsupervised process (Frémaux and Gerstner, 2015), i.e., it is not known whether feature creation is influenced by rewarding or aversive signals. So, in a general theory, it might be valid to consider feature learning as a separate, unsupervised, module. This approach could lead to a conceptually simple framework of agent behavior and feature detection simultaneously. Concretely, the approach could enable the results in this paper to be linked to various feature detector algorithms.

The challenge of new features In this paper we have assumed that the set of features is fixed at the beginning of the learning process. This could be suitable for many applications, as there is no fixed limit on how many features there are, as long as there are finitely many. But it seems intriguing to introduce new features while the agent is performing the task. In the technical approach of this paper, however, a newly inserted feature likely proposes wrong actions if we would initially associate all actions to the feature. In general we still insist that aversive signals are avoided, and therefore the wrong actions need to be unlearned as soon as possible.

A way to soften the introduction of new features, could be to reintroduce reward into the framework. Concretely, a feature f may only propose an action a if the pair (f, a) has been observed to be correlated to reward, either directly, or transitively by means of eligibility traces (Sutton and Barto, 1998). This idea introduces a threshold for proposing actions. Of course any feature-action pairs introduced in this way could still lead to aversive signals. For example, there could be spurious features (e.g. features that randomly appear) to which no actions should be linked, or perhaps the rewarding signals contradict the aversive signals, or some actions that give reward could also give aversive signals (as in the example of the Introduction). To resolve priority issues, one could view avoidance learning as having the highest precedence, where reward is used as a softer ranking mechanism on the allowed actions.

Possibly, an agent that keeps learning new features will keep making mistakes. How to cope with new features therefore seems a relevant question. The answers could perhaps also help to understand animal behavior and consciousness. Thereto one could consider other notions of success than the avoidance of aversive signals investigated in this paper.

References

Chatterjee, K. and Chmelík, M. (2015). POMDPs under probabilistic semantics. *Artificial Intelligence*, 221:46 – 72.

Francez, N. (1986). *Fairness*. Springer-Verlag New York, Inc.

Frémaux, N. and Gerstner, W. (2015). Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. *Frontiers in Neural Circuits*.

Frémaux, N., Sprekeler, H., and Gerstner, W. (2013). Reinforcement learning using a continuous time actor-critic framework with spiking neurons. *PLoS Computational Biology*, 9(4):e1003024.

García, J. and Fernández, F. (2015). A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16:1437–1480.

Geva-Sagiv, M., Las, L., Yovel, Y., and Ulanovsky, N. (2015). Spatial cognition in bats and rats: from sensory acquisition to multiscale maps and navigation. *Nature Reviews Neuroscience*, 16(4):94–108.

Heger, M. (1994). Consideration of risk in reinforcement learning. In *Proceedings of the Eleventh International Conference on Machine Learning*, pages 105–111.

Jaakkola, T., Jordan, M., and Singh, S. (1994). On the convergence of stochastic iterative dynamic programming algorithms. *Neural Computation*, 6(6).

Lusena, C., Goldsmith, J., and Mundhenk, M. (2001). Nonapproximability results for partially observable markov decision processes. *Journal of Artificial Intelligence Research*, 14:83–103.

Pfeifer, R. and Scheier, C. (1999). *Understanding Intelligence*. The MIT Press.

Potjans, W., Diesmann, M., and Morrison, A. (2011). An imperfect dopaminergic error signal can drive temporal-difference learning. *PLoS Computational Biology*, 7(5):e1001133.

Roy, N., Gordon, G., and Thrun, S. (2005). Finding approximate POMDP solutions through belief compression. *Journal of Artificial Intelligence Research*, 23:1 – 40.

Sutton, R. and Barto, A. (1998). *Reinforcement Learning, An Introduction*. The MIT Press.

van Otterlo, M. (2009). *The logic of adaptive behavior*. IOS Press.

Watkins, C. (1989). Learning from delayed rewards. *PhD thesis, Cambridge University*.

Watkins, C. and Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3–4):279–292.